

**SYSTEM AND METHOD FOR IMPROVED UTILIZATION OF
BANDWIDTH IN A COMPUTER SYSTEM SERVING MULTIPLE USERS**

Inventors:

Robert S. Horen

Sanjay Singal

RELATED APPLICATIONS

This application claims priority to co-pending United States Provisional Patent Application Serial No. 60/221,593, filed July 28, 2000, with the United States Patent and Trademark Office, which is incorporated herein by reference.

FIELD OF THE INVENTION

The present invention relates broadly to distributed computer systems. More specifically, the present invention relates to organizing content into asset groups to optimize delivery to users in a distributed computer system.

BACKGROUND

In a computer system serving media assets to multiple users, such as a server farm having significant mass storage in the form of magnetic and optical disks and providing content to users of client computers over a global computer network, a value indicating a guaranteed number of plays is an attribute associated with each media asset. This attribute is used to reserve storage bandwidth so that at any given time, a certain number of simultaneous playouts are available for the requested media asset. For file systems containing the media assets and as referred to herein, storage bandwidth refers

to the transmission bandwidth available for transfer of media assets stored on the server computer over a communications medium, such as a network connection. For example, the storage bandwidth for one media asset doubles when it two media assets of the same size are played out simultaneously. A media asset can for example be audio, text, graphics, image, symbol, video, or other information item or token or communication item or method by which information such as audio or visual content is conveyed to a user. As referred to herein, "playout" refers to the streaming or other transfer of the media asset from a server computer to a client computer on which a user views or listens to the media asset. Additionally, the guaranteed number of plays attribute is used to determine the number of copies of the asset required to satisfy the guaranteed number of plays. However, there are limitations associated with utilization of this attribute. As shown in FIG. 1, a significant amount of storage space is wasted with respect to the storage bandwidth utilized for playouts of the ten media assets.

FIG. 1 illustrates utilization of storage space on a server versus storage bandwidth as utilized in conventional media asset delivery systems. In FIG. 1, a user desires a maximum of ten playouts from a set of ten 1Mbps assets, each of which occupies ten MB of storage space. These playouts could all be from the same asset or from any combination thereof, although the total number of playouts is less than or equal to the ten playouts desired by the user. Typical implementations install each of these assets with ten guaranteed playouts. If these assets are all placed on a single file system with a bandwidth capacity of 100Mbps and space capacity of 1GB, then the entire file system bandwidth is consumed and the file system is no longer usable for further asset installations even though only 100MB of disk space has been used. This is wasteful in terms of file system bandwidth and space.

SUMMARY

The present invention solves the problems discussed above by providing a method, computer program product and server computer system for use in a client server computer architecture. The server sends media assets over a computer network to a client computer and maintains a file system organized into a plurality of asset groups, each asset group comprising a plurality of media assets, wherein the media

assets share storage medium bandwidth and storage space on the server computer that is reserved for the asset group to which the plurality of media assets belong.

Attributes are associated with each asset group, and can include a value indicating the number of maximum simultaneous playouts for the media assets within the asset group, and the maximum bit rate at which any single media asset within the asset group can be played out, the guaranteed possible number of playouts from each asset belonging to the asset group.

An asset group placement policy module is provided that evaluates the attributes of the asset group to determine optimal placement of the asset group within the file system of the server computer system, avoiding replication of assets and spreading the asset group across multiple file systems wherever possible.

A media asset placement policy module is provided that evaluates asset bandwidth to determine the optimal placement for the asset and available resources and use this evaluation to distribute the media assets within the asset groups.

By organizing media assets into asset groups of the present invention, resources such as storage bandwidth and storage space are conserved, thus allowing a server computer system to play a greater number of media assets to clients than conventional systems of similar storage capacity and storage bandwidth.

These and many other attendant advantages of the present invention will be understood upon reading the following detailed description in conjunction with the drawings.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 is a diagram of utilization of bandwidth and storage space as used in prior art systems;

FIG. 2 is a diagram of utilization of bandwidth and storage space as used in accordance with an embodiment of the present invention;

FIG. 3 is a block diagram illustrating an embodiment of the present invention as implemented in a client-server computer system implemented over a computer network;

FIG. 4 is a block diagram illustrating the major components of a computer that may be used in connection with the client server system according to an embodiment of the invention;

FIG. 5 is a table illustrating demand for media assets in a distance learning application according to an embodiment of the present invention;

FIG. 6 is an illustration of an embodiment of the present invention as used to partition media assets between groups of users according to an embodiment of the invention; and

FIG. 7 is a flow chart diagram illustrating the sequence of operations executed to perform a method of an embodiment of the present invention.

DETAILED DESCRIPTION

Directing attention to FIG. 2, the present invention avoids the shortcomings of the problems discussed above and illustrated in FIG.1 by organizing media assets into asset groups. As referred to herein, an asset group is a set of media assets (static image, video, audio, text, or combinations thereof) that can be played on a computer and which share a certain amount of system resources, such as storage space and storage bandwidth. FIG. 2 illustrates the manner in which asset groups can be used so that the storage bandwidth is more effectively used in comparison to the rapid growth or consumption of bandwidth on the conventional system and method of FIG. 1.

Directing attention to FIG. 3, embodiments of the present invention utilize a client server computer architecture 100 having a server 102 connected or capable of being connected to one or more client computers 104-1, 104-2,..., 104- n , where n is the number of consumers that utilize client computer 104 to receive content from server 102 via computer network 106. In the preferred embodiment, computer network 106 is a global computer network, such as the Internet, but may also be a wide area network, such as used within a company having multiple facilities located at different geographic sites, or smaller computer networks, such as computer networks used on college campuses, local area networks, and the like.

Server 102 incorporates connection server 108, a module that resides in memory and coordinates connections between the server 102 and clients 104. Once a connection is made, the server 102 then directs the distribution of media assets 110 in response to

user requests for delivery. Media assets 110 can be stored in memory and/or mass storage devices, depending on user demand and storage requirements of the individual assets. The media assets 110 are organized within file systems 112 that are configured to store the assets in asset groups 114. An asset group 114 is normally administered by a single asset group administrator 116. In an embodiment of the present invention, each asset belongs to only one asset group, thus avoiding replication of assets and more efficiently using storage space and storage bandwidth. Asset group 114 can contain other asset groups as well as other assets, much as a file directory can contain files as well as subdirectories of files. Different assets within asset group 114 can have varied bit rates. The following attributes can be associated with asset group 114: Maximum Simultaneous Plays for Asset Group (118), Maximum bit rate (120), Default Guaranteed Possible Plays (DGPP) (122), Guaranteed Possible Playouts (124) (GPP), and Resource Quota (126). Each of these attributes are described in further detail immediately below.

Maximum Simultaneous Playouts for Asset Group (118) is an attribute that indicates the maximum simultaneous playouts possible from asset group 114. Sufficient file system bandwidth is reserved at asset group creation time in order to satisfy this value.

Maximum bit rate (120) is an attribute that describes the maximum bit rate of any single asset installed within asset group 114.

Default Guaranteed Possible Playouts (DGPP) (122) is an attribute that indicates the guaranteed number of playouts possible from each asset 110 within the asset group 114 assuming no other asset is being played out at the same time. Internally, the storage manager 123 will create enough copies of assets within asset group 114 at install time to satisfy the value designated by the DGPP 122. This value can be overridden for each asset 110.

Guaranteed Possible Playouts (124) (GPP): note that if A_1, A_2, \dots, A_n are the assets within an asset group, $GPP_1, GPP_2, \dots, GPP_n$ are the GPP attributes for these assets respectively, and b_1, b_2, \dots, b_n are the bit rates of the assets respectively, and BWAG is the bandwidth quota for asset group 114, then the following inequality holds: $GPP_i * b_i \leq BWAG$ for $i=1,2,\dots,n$.

Resource Quota (126) is an attribute that indicates the file system resources (storage space and storage bandwidth) reserved for asset group 114. All assets within asset group 114 share the storage bandwidth, and are not played out using any asset group's bandwidth. The bandwidth component of resource quota 126 is calculated by multiplying the value of the Maximum Simultaneous Playouts for Asset Group attribute 118 by the value of the Maximum Bit Rate attribute 120. Storage space computation is slightly more involved and depends on the maximum bit rate, duration of playing assets installed within asset group 114, the DGPP value for asset group 114 and also the availability of storage space and bandwidth on currently configured file system 112. The assets 110 installed within asset group 114 may not exceed the resource quota 126 in order to preserve the reservation of resources for each asset group.

With the introduction of asset groups, two levels of placement policies are implemented. Asset group placement policy module 128 performs asset group placement based on attributes which determine the resource quota 126 for the asset group 114, and media asset placement policy module 130 performs media asset placement within the placed asset group. In order to better utilize the resource quota, avoid fragmentation of disk resources and avoid asset replication where possible, asset groups are placed in a manner that avoid distributing the quota across multiple file systems as much as possible, such as on the same disk array or physical storage device, or within the group storage devices that make up a single file system. Policy module 130 evaluates asset bandwidth to determine the optimal placement for the asset and available resources and use this evaluation to distribute the media assets within the asset groups 114. Since storage space and storage bandwidth reservation have already been performed for the asset group 114 by policy module 128 prior to asset installation, policy module 130 restricts the placement domain of the media asset to the asset group distribution of storage space and storage bandwidth.

FIG. 4 illustrates in block diagram form the major components included in a computer embodying either server 102 or client 104. Computer 200 incorporates a processor 202 utilizing a central processing unit (CPU) and supporting integrated circuitry. In the preferred embodiment, work stations such as Sun Ultra computers available from Sun Microsystems can be used as server 102. Personal computers such

as available from Dell Corporation may be used for client computers 104. However, in general, any type of computer may be used for the server 102 and any type of computer may be used for client 104. Memory 204 may include RAM and NVRAM such as flash memory, to facilitate storage of software modules executed by processor 202, and file systems containing media assets. Also included in computer 200 are keyboard 206, pointing device 208, and monitor 210, which allow a user to interact with computer 200 during execution of software programs. Mass storage devices such as disk drive 212 and CD ROM 214 may also be in computer 200 to provide storage for computer programs, associated files, and media assets. In an embodiment, database products available from Oracle Corp. may be utilized in connection with file systems as a database and database server. Computer 200 communicates with other computers via communication connection 216 and communication line 218 to allow the computer 200 to be operated remotely, or utilize files stored at different locations. Communication connection 216 can be a modem, network interface card, or other device that enables a computer to communicate with other computers. Communication line 218 can be a telephone line or cable, or any medium capable of transferring data between computers. In alternative embodiments, communication connection 216 can be a wireless communication medium, thus eliminating the need for communication line 218. The components described above may be operatively connected by a communications bus 220.

Embodiments of the present invention are useful in a variety of applications where multiple media assets share a finite amount of bandwidth. Two scenarios where asset groups of the present invention provide improved file system resource utilization are described below. These examples are provided to illustrate the benefits of the present invention and are just two of many applications for which the present invention is suitable.

In an embodiment of the present invention, consider the situation of a news acquisition system as used by a major news provider such as CNN. Such a system provides asset playout access to a maximum of approximately 300 journalists. Assets are captured using a set of 40 encoders. Assets being captured (and stored in asset groups) comprise breaking news items. For illustrative purposes it is assumed that during a typical network load ten viewers access a given news item simultaneously.

There can be multiple stories (up to 40 in this scenario, limited by the number of encoders) being captured at any given time. Additionally, there can be numerous old stories stored on the system.

According to one embodiment, an asset group is defined with the following attributes: (a) Maximum number of simultaneous playouts for asset group = 300, (b) Maximum bit rate of assets installed = 1.5 Mbps, (c) Default Guaranteed Possible Playouts = 10. Given these attributes and values, the bandwidth component of the resource quota for the asset group is determined to be 450 Mbps. In the case of three available file systems: F1, having 300 Mbps of storage bandwidth and 60 GB of storage space; F2 having 150 Mbps of storage bandwidth and 30 GB of storage space; and F3, having 120 Mbps of storage bandwidth and 60GB of storage space, an asset group can be created by the asset group placement policy 128 with the following distribution: ((F1, 250Mbps, 50GB), (F2, 100Mbps, 20GB), (F3, 100Mbps, 50GB)). Since the DGPP is 10, the maximum bit rate allowed for any single asset is 15Mbps. This means that asset replication can be completely avoided. If the length of each news story is 90 minutes (which equates to about 1.0125GB) this asset group can accommodate a total of about 175 hours of news stories or 117 news stories.

Contrast the embodiment of the present invention described immediately above with an implementation of a conventional system. If the media assets were installed individually on this system with guaranteed playouts = 10, the total number of assets installed on the entire system could not have exceeded 570 Mbps, which is the combined bit rates for file systems F1, F2, and F3, divided by 15 Mbps, which is the maximum bit rate for any single asset. 570 Mbps divided by 15 Mbps yields 38 news stories. By forming an asset group, 117 news stories, nearly triple the amount of media assets allowed by conventional systems, may be accommodated using only a fraction of the file system resources.

In another embodiment of the present invention, consider the example of a university which would like to offer lecture notes for three courses to be viewed by two kinds of students: (a) students on campus who have access to the high-speed campus backbone and can view assets at 1.5 Mbps, and (b) students accessing the lecture notes over the Internet and thus limited to about 200kbps. Further, the university provides lecture notes for the past ten weeks on a sliding window basis. Lecture notes for the

current week are likely to be the most in demand, while those for the previous weeks are less in demand. The demand table to ten weeks is shown in the table in FIG. 5. The first column indicates the course name, the second the bit rate of assets available for that course (the "i" suffix for the course name indicates an Internet offering). The next ten columns indicate the maximum possible number of students who can access the course offerings for that particular week. For example, during the first week a maximum of fifty students can access the lecture notes for MA101, and during the seventh week a maximum of twenty students can access the lecture notes for CH101i. In addition, the number in parentheses "()" next to the week number indicates the maximum number of simultaneous accesses for each individual lecture in the week. There are five sets of lecture notes stored for each week in each category. This example is ideal again for the illustration of the effectiveness of asset groups. The lecture notes administrator can create separate asset groups for the lecture notes of each course offering each week for each kind of bit rate. In this example, this means that there is a total of 60 asset groups, determined by each course having an asset group for the lecture notes for each week of the course. Half of these asset groups are assumed to have a maximum bit rate of 1.5Mbps, the other half 200kbps. The asset group for the lecture notes for the campus-access version of CS101 for week one would have the following attributes: (a) Maximum Simultaneous Playouts for Asset Group = 50, (b) DGPP = 20, and (c) Maximum Bit Rate for Assets Installed in Asset Group = 1.5 Mbps. Let us further assume that each lecture is 60 minutes in duration.

System bandwidth and space requirements are now addressed. For this exemplary computation it is assumed that each lecture is 60 minutes in duration. Assuming no replication is needed (for simplicity, but not limiting the scope of the invention), the file system bandwidth and space requirement for this kind of a setup are as follows:

Week 1:

$$50 * 1.5 \text{ Mbps} * 3 \text{ courses} + 50 * 0.2 \text{ Mbps} * 3 \text{ courses} = \mathbf{255 \text{ Mbps bandwidth, and}}$$

$$(5 * 60 * 60 * 1.5 \text{ Mbps} * 3 \text{ courses} + 5 * 60 * 60 * .2 \text{ Mbps} * 3 \text{ courses}) / 8 = 11.475 \text{ GB}$$

If on the other hand, asset groups were not used and individual assets were installed, the following would have been the resource requirements:

Week 1:

$$5 * 20 * 1.5 * 3 + 5 * 20 * 0.2 * 3 = 510 \text{ Mbps bandwidth.}$$

The storage space is the same:

$$(5 * 60 * 60 * 1.5 \text{ Mbps} * 3 \text{ courses} + 5 * 60 * 60 * .2 \text{ Mbps} * 3 \text{ courses}) / 8 = 11.475 \text{ GB}$$

5

Thus, the present invention reduces the bandwidth requirement by 50%, from 510 Mbps to 255 Mbps. Note that the space calculation here assumes that no space is wasted due to unavailability of bandwidth on certain disks.

Directing attention to FIG. 6, an embodiment of the present invention can also be used to partition file system resources between groups of users. In a server having a file system bandwidth of for example 100 Mbps and 10GB of storage for media assets, asset groups can be used to partition file system resources, with each asset group designated for access to a particular user group. In this example, asset group 310 is assigned to a first user group and has 2 GB of storage and a bandwidth of 30Mbps, asset group 312 is assigned to a second user group and has 2 GB of storage and a bandwidth of 30Mbps, and asset group 314 is assigned to a third user group and has 6 GB of storage and a bandwidth of 40Mbps. Using some primitives, such as authorization, authentication and accounting products, available from Portal Software in Cupertino, California, it is possible to provide tracking, billing, monitoring, and other services for all asset groups.

Directing attention to FIG. 7, the method 400 of the present invention to create asset group 114 is performed through the following sequence of steps. At step 402, asset group 114 is created, including attributes Maximum Simultaneous Playouts for Asset Group (118), Maximum Bit Rate of the Assets Installed in Asset Group (120), Default Guaranteed Possible Playouts (122), and Resource Quota (126). Creating the asset group 114 can be performed by defining a data structure that contains a list of pointers that contain values indicating the storage locations of the media assets stored in the asset group, as well as pointers to values of the attributes or simply the attribute values themselves rather than pointers. At step 404, the storage bandwidth component of the Resource Quota is calculated by multiplying together the Maximum Simultaneous Playouts for Asset Group (118), Maximum Bit Rate (120), Default Guaranteed Possible Playouts (DGPP) (122). At step 406, the asset group 114 is

10
15
20
25
30

assigned to a file system 112 maintained on server 102. At step 410, media assets (including media asset 110) are stored in accordance with the asset group 114 and made available to client computer 104.

Having disclosed exemplary embodiments and the best mode, modifications and variations may be made to the disclosed embodiments while remaining within the scope of the present invention as defined by the following claims.

1032313

09916666-072701
10/22/20" 55331660